

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:	§	Examiner: Joo, Joshua
Bernard A. Traversat, et al.	§	
	§	
	§	Group Art Unit: 2154
Serial No. 10/055,650	§	
	§	Atty. Dkt. No.: 5681-07400
Filed: January 22, 2002	§	
	§	
For: Reliable Peer-to-Peer	§	
Connections	§	
	§	

APPEAL BRIEF

Mail Stop Appeal Brief - Patents

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal filed June 23, 2008, Appellants present this Appeal Brief. Appellants respectfully request that the Board of Patent Appeals and Interferences consider this appeal.

I. REAL PARTY IN INTEREST

As evidenced by the assignment recorded at Reel/Frame 012890/0833, the subject application is owned by Sun Microsystems, Inc., a corporation organized and existing under and by virtue of the laws of the State of Delaware, and now having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054.

II. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences or judicial proceedings are known which would be related to, directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-64 are pending in the application and stand finally rejected. The rejection of claims 1-64 is being appealed. A copy of claims 1-64 is included in the Claims Appendix herein below.

IV. STATUS OF AMENDMENTS

No amendments have been submitted subsequent to the final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a peer computing system, including a plurality of peer nodes coupled to a network. (*See, e.g., FIGs. 1A and 1B, peer devices 104; and p. 4, lines 21-26.*)

Each of the peer nodes includes one or more network interfaces, and each network interface is configured to communicate over the network in accordance with at least one of one or more network transport protocols. (*See, e.g., FIG. 30, peers 200; p. 32, lines 13-21; and p. 43, lines 13-18.*)

The peer nodes are configured to implement a peer-to-peer environment on the network according to a peer-to-peer platform that includes one or more peer-to-peer platform protocols for enabling the peer nodes to discover each other, communicate with each other, and share content in the peer-to-peer environment, where discovering each other includes obtaining an address for each discovered peer node. (*See, e.g., FIG. 2, P2P platform core 120; p. 18, lines 15-23; p. 19, lines 4-17; p. 20, lines 4-28; and p. 95, lines 1-15.*)

One of the peer nodes is configured to establish a communications channel between a network interface of the peer node and a network interface of another of the peer nodes. (*See, e.g., FIG. 30, peers 200; p. 32, lines 13-21; and p. 43, lines 13-18.*)

The peer node is also configured to transmit messages to the other peer node over the communications channel, receive acknowledgement that one or more of the transmitted messages have been received by the other peer node; and retransmit messages not acknowledged as received by the other peer node to the other peer node on the communications channel. (*See, e.g., FIG. 33A, communications channel 416; and p. 51, line 24 – p. 52, line 25.*)

Establishing a communications channel, transmitting messages, receiving

acknowledgment, and retransmitting messages are performed according to at least one of the one or more peer-to-peer platform protocols, which are distinct from the network transport protocols. (See, e.g., p. 22, lines 18-28; and page 78, lines 24-27.)

Independent claim 25 is directed to a method for providing reliable connections between peer nodes coupled to a peer-to-peer network. (See, e.g., FIGs. 1A and 1B, peer devices 104; and p. 4, lines 21-26.)

The method includes a plurality of peer nodes coupled to the network implementing a peer-to-peer environment on the network according to a peer-to-peer platform. The peer-to-peer platform includes one or more peer-to-peer platform protocols for enabling the peer nodes to discover each other, communicate with each other, and share content in the peer-to-peer environment, wherein discovering each other includes obtaining an address for each discovered peer node. (See, e.g., FIG. 2, P2P platform core 120; p. 18, lines 15-23; p. 19, lines 4-17; p. 20, lines 4-28; and p. 95, lines 1-15.)

The method also includes establishing a communications channel between a network interface of one of the peer nodes and a network interface of another of the peer nodes. (See, e.g., FIG. 30, peers 200; p. 32, lines 13-21; and p. 43, lines 13-18.)

The method also includes the peer node transmitting messages to the other peer node over the communications channel, receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, and retransmitting messages not acknowledged as received by the other peer node to the other peer node on the communications channel. (See, e.g., FIG. 33A, communications channel 416; and p. 51, line 24 – p. 52, line 25.)

Establishing a communications channel, transmitting messages, receiving acknowledgement, and retransmitting messages are performed according to at least one

of the peer-to-peer platform protocols, which is distinct from any underlying network transport protocols. (*See, e.g., p. 22, lines 18-28; and page 78, lines 24-27.*)

Independent claim 45 is directed to a computer-readable storage medium, storing software instructions executable on a peer node to implement the peer node establishing a communications channel between a network interface of the peer node and a network interface of another peer node of a plurality of peer nodes. (*See, e.g., FIG. 30, peers 200; p. 32, lines 13-21; p. 43, lines 13-18; and p. 137, lines 10-12.*)

The plurality of peer nodes is coupled to a network and implements a peer-to-peer environment on the network. (*See, e.g., FIGs. 1A and 1B, peer devices 104; and p. 4, lines 21-26.*) The peer-to-peer environment is implemented according to a peer-to-peer platform that includes one or more peer-to-peer platform protocols for enabling the peer nodes to discover each other, communicate with each other, and share content in the peer-to-peer environment, wherein discovering each other includes obtaining an address for each discovered peer node. (*See, e.g., FIG. 2, P2P platform core 120; p. 18, lines 15-23; p. 19, lines 4-17; p. 20, lines 4-28; and p. 95, lines 1-15.*)

The software instructions are also executable to implement the peer node transmitting messages to the other peer node over the communications channel, receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, and retransmitting messages not acknowledged as received by the other peer node to the other peer node on the communications channel. (*See, e.g., FIG. 33A, communications channel 416; and p. 51, line 24 – p. 52, line 25.*)

Establishing a communications channel, transmitting messages, receiving acknowledgement, and retransmitting messages are performed according to at least one of the peer-to-peer platform protocols, which is distinct from any underlying network transport protocols. (*See, e.g., p. 22, lines 18-28; and page 78, lines 24-27.*)

The summary above describes various examples and embodiments of the claimed subject matter; however, the claims are not necessarily limited to any of these examples and embodiments. The claims should be interpreted based on the wording of the respective claims.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1-3, 5-7, 11-15, 18, 21, 22, 25-27, 29-31, 35-40, 43, 45-47, 49-51, 55-60 and 63 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Davis et al. (U.S. Patent 6,105,064) (hereinafter “Davis”) in view of Dreke et al. (U.S. Publication 2002/0035594) (hereinafter “Dreke”) and Black et al. (U.S. Patent 5,878,056) (hereinafter “Black”).

2. Claims 4, 8-10, 28, 32-34, 48 and 52-54 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Davis, Dreke and Black in view of Barker et al. (U.S. Patent 5,931,916) (hereinafter “Barker”).

3. Claims 16 and 17 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Davis, Dreke and Black in view of Ivanoff (U.S. Patent 5,517,622) (hereinafter “Ivanoff”).

4. Claims 19, 20, 41, 42, 61 and 62 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Davis, Dreke and Black in view of Antur et al. (U.S. Patent 6,212,558) (hereinafter “Antur”).

5. Claims 23, 24, 44 and 64 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Davis, Dreke and Black in view of Zhu et al. (U.S. Patent 5,768,557) (hereinafter “Zhu”).

VII. ARGUMENT

First ground of rejection:

The Examiner rejected claims 1-3, 5-7, 11-15, 18, 21, 22, 25-27, 29-31, 35-40, 43, 45-47, 49-51, 55-60 and 63 under 35 U.S.C. § 103(a) as being unpatentable over Davis et al. (U.S. Patent 6,105,064) (hereinafter “Davis”) in view of Dreke et al. (U.S. Publication 2002/0035594) (hereinafter “Dreke”) and Black et al. (U.S. Patent 5,878,056) (hereinafter “Black”). Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Claims 1, 2, 11, 15, 18, 21, 22, 25, 26, 35, 39, 40, 43, 45, 46, 55, 59, 60, and 63:

1. **The cited art clearly fails to teach or suggest *wherein said establishing, said transmitting, said receiving, and said retransmitting are performed according to at least one of the one or more peer-to-peer platform protocols and wherein said peer-to-peer platform protocols are distinct from the at least one network transport protocols, as recited in claim 1.***

As discussed in Appellants’ previous Responses, the invention disclosed in Davis is directed to network transport layers and network transport protocols, as described in the Field of Invention, “The present invention relates to communication over a computer network, and more particularly to a approach which employs dynamic window sizing, packet metering, and other techniques to provide an efficient and reliable network transport layer” (emphasis added). Various examples of such network transport protocols are also described in Davis, including ACP (an embodiment of Davis’ claimed invention), TCP (Transmission Control Protocol) and UDP (User Datagram Protocol) on UNIX systems, protocols TP0 through TP4 of the OSI model, SPX (Sequenced Packet Exchange) and IPX (Internet Packet Exchange) in Novell NetWare systems (SPX, IPX, NOVELL, and NETWARE are trademarks of Novell, Inc.), and other protocols. Appellants note that many of the Examiner’s cited passages, including those in columns 5

and 9 describe network transport protocols, as stated in column 5, lines 40-44, “**More specifically, the present invention provides several protocols for controlling data packets at the transport layer or other packet transmission layer**” (emphasis added). **In other words, the protocols taught by Davis are explicitly described as being implemented using network transport protocols.**

The Examiner submitted that Black teaches implementing a messaging system that is independent of transport protocols, in column 10, lines 63-67. This passage actually states, “The message format and the safe movement protocol are transport layer independent so that MCAs can support different transport protocols on different channels” (emphasis added). Appellants assert that there is nothing in Black that describes the “message format” and “safe movement protocol” as peer-to-peer platform protocols, or anything about a peer computing system, at all. Therefore, having a “message format” and a “safe movement protocol” that are transport layer independent does not teach or suggest that peer-to-peer platform protocols, such as the specific peer-to-peer platform protocols recited in claim 1, should be transport layer independent. **None of the cited references include such protocols, and the system of Davis is specifically directed to network transport protocols (e.g., ACP).**

In the Response to Arguments section of the Final Action mailed March 21, 2008, the Examiner disagreed, submitting, “The said establishing, said transmitting, said receiving, and said retransmitting between the peers are considered as the peer-to-peer platform protocols. Davis does not specifically teach that the peer-to-peer platform protocols are distinct from the at least one network transport protocols. Black teaches of a system for a messaging protocol between devices that is independent, i.e., distinct from a network transport protocol.” **Appellants again asserted that none of the cited references teach the peer-to-peer platform protocols of claim 1. Therefore, the combination of references cannot teach that such protocols are (or would benefit from being) distinct from network transport protocols.**

2. **The cited art clearly fails to teach or suggest *wherein the plurality of peer nodes is configured to implement a peer-to-peer environment on the network according to a peer-to-peer platform comprising one or more peer-to-peer platform protocols, as recited in claim 1.***

The Examiner cited Davis column 8, lines 21-24 as teaching the peer-to-peer platform protocols of claim 1. However, as noted in Appellants' previous Response, this passage merely states that a given computer "may also function as a peer in a peer-to-peer network" without describing any of the specific limitations of a peer-to-peer platform recited in the claims or mentioning one or more peer-to-peer platform protocols, as in claim 1. **Appellants again assert that a computer may "function as a peer in a peer-to-peer network" without necessarily including a peer-to-peer platform comprising any of the specific peer-to-peer platform protocols recited in claim 1, and without meeting the additional limitations recited therein regarding their distinction from transport protocols.** As discussed at length in Appellants' previous Responses, Appellants assert that none of the cited passages, or anything else in Davis, describes peer-to-peer platform protocols of a peer-to-peer platform, meeting the limitations of Appellants' claims. Davis is concerned with dynamically adjusting the propagation rate of packets between known sending and receiving nodes. Davis teaches the use of protocols such as TCP, UDP, SPX, IP, IPX and ATM, all of which are explicitly described in Davis as being network transport protocols.

Davis, Dreke, and Black do not describe the particular peer-to-peer platform protocols recited in claim 1. For example, these references fail to teach or suggest any peer-to-peer platform protocols for enabling peers to discover each other, *wherein to discover comprises obtaining an address for each discovered peer node*. The Examiner's cited passage in Dreke, for example, describes three client computers establishing a connection through an Internet Presence Information Server (IPIS), "In 301, Peer A first transmits to IPIS 4 the following information: his/her newly assigned network (Internet Provider (IP)) address; a list of peers whose Internet presence are of interest to Peer A; and a request for a list of peers who are interested in the Internet presence of Peer A. In

this example, the list transmitted by Peer A includes Peer B and Peer C. In 302, IPIS 4 responds to Peer A's list by transmitting a list including the last known address, such as an IP addresses for Peer B and Peer C even though the IP address for Peer B is out of date. During 302, IPIS 4 also responds to Peer A's request for a list of peers interested in Peer A's presence with a message indicating no peers are currently interested in his/her presence. Once IPIS 4 transmits these lists to Peer A, Peer A will no longer communicate with IPIS 4 during this network session.” This clearly does not describe a peer-to-peer platform protocol for enabling peers to discover each other, as recited in 1. Instead this describes a mechanism to establish a connection between two known peers, in which locating the peers and establishing the connection are managed by the IPIS server. **By definition**, the centralized IPIS sever mechanism of Dreke does not involve a peer-to-peer platform protocol for enabling peer nodes to discover each other. Appellants further noted that there are many ways (such as that described in Dreke) that a device may obtain an address for another device that do not involve a peer-to-peer platform protocol, much less one that is distinct from network transport protocols. The cited references, whether considered singly or combination, clearly do not describe the peer-to-peer platform protocol for enabling peer nodes to discover each other, according to the limitations of claim 1.

3. The Examiner has not provided a proper reason to combine the references.

In the Response to Arguments section of the Final Action mailed March 21, 2008, the Examiner submitted that it would have been obvious to one of ordinary skill in the art to implement the peer-to-peer platform protocols as taught by Davis as protocol independent of transport protocols as taught by Black, which would improve Davis' teachings by allowing the peer-to-peer protocol including said steps as taught by Davis to be implemented regardless of the transport protocols and provide reliable data transmission for different transport protocols. The Examiner's stated reason for combining the references is unsupported by the cited art. As discussed above, Davis is directed to techniques to provide an efficient and reliable network transport layer, i.e., to

network transport protocols. Therefore, it does not make sense to implement the teachings of Davis independent of such transport protocols, nor is it clear if or how this could be accomplished. The protocols taught by Davis are explicitly described as being implemented using network transport protocols. Therefore, the modification suggested by the Examiner would change the principle of operation of the system of Davis. Appellants remind the Examiner that “If the proposed modification or combination of the prior art would change the principle of operation of the prior art invention being modified, then the teachings of the references are not sufficient to render the claims *prima facie* obvious.” *In re Ratti*, 270 F.2d 810, 123 USPQ 349 (CCPA 1959). **Furthermore, since Davis does not teach the peer-to-peer protocols of Appellants’ claim, the combination of the references would still not result in Appellants’ claimed invention.**

In the Advisory Action mailed June 9, 2008, the Examiner again asserts that it would have been obvious to one of ordinary skill in the art to modify Davis’ teachings for the protocols to be implemented as a messaging protocol independent of the network transport protocol as taught by Black, which would allow Davis’ protocols to operate on different transport protocols. **Appellants again assert that this argument is nonsensical, since Davis’ teachings are directed to the very network transport protocols themselves.**

To establish a *prima facie* obviousness of a claimed invention, all claim limitations must be taught or suggested by the prior art. *In re Royka*, 490 F.2d 981, 180 U.S.P.Q. 580 (C.C.P.A. 1974), MPEP 2143.03. As shown above, there is no doubt that the cited art does not teach or suggest all limitations of the claim 1, nor has the Examiner provided a sufficient reason to combine the references.

For at least the reasons above, the rejection of claim 1 is not supported by the cited art and removal thereof is respectfully requested. Independent claims 25 and 45 include limitations similar to the above-referenced limitations of claim 1 and were

rejected for similar reasons. Therefore, the arguments presented above apply with equal force to these claims, as well.

Dependent claims 3, 5, 6, 27, 29, 30, 47, 49, and 50:

1. **The cited art clearly fails to teach or suggest *wherein the other peer node is configured to receive the transmitted messages, and after receiving M messages, transmit the acknowledgement to the peer node indicating that the M messages have been received, wherein M is a positive integer less than or equal to N , as recited in claim 3.***

The Examiner submits that Davis teaches these limitations, citing column 30, lines 66-67 and column 59, lines 34-35. The first of these passages describes that a send window size may be increased by the number of packets acknowledged by a received ACK. The second passage describes only that a receiving endnode may not necessarily send one acknowledgement packet for each received packet. Neither of these passages describes transmission of an acknowledgement indicating that M messages have been received. Instead, they describe only that an acknowledgement may not be sent after each packet is received. In other passages in Davis, it appears that only the most recently received packet, or the most recently received packet in sequential order, is acknowledged, and nothing describes that it includes an indication of a plurality of messages, as required by claim 3.

For at least the reasons above, the rejection of claim 3 is unsupported by the cited art and removal thereof is respectfully requested.

Claims 27 and 47 include limitations similar to those discussed above regarding claim 3. Therefore, the arguments presented above apply with equal force to this claim, as well.

Dependent claims 7, 31, and 51:

1. The cited art clearly fails to teach or suggest *wherein the shifted window includes one or more messages previously transmitted to the other peer node and one or more messages not previously transmitted to the other peer node*, as recited in claim 7.

The Examiner submits that Davis teaches these limitations in column 30, lines 1-8 and column 29, lines 51-60. Appellants note that column 30 does not include lines 1-8. Therefore it is unclear what the Examiner means to reference. Appellants assert that there is nothing in the Examiner's second passage that describes the shifted window including one or more messages previously transmitted and one or more messages not previously transmitted, as the Examiner suggests.

For at least the reasons above, the rejection of claim 7 is unsupported by the cited art and removal thereof is respectfully requested.

Claims 31 and 51 include limitations similar to those discussed above regarding claim 7. Therefore, the arguments presented above apply with equal force to this claim, as well.

Dependent claims 12-14, 36-38, and 56-58:

1. The cited art clearly fails to teach or suggest *wherein the peer node and the other peer node are further configured to: monitor reception and retransmission of the messages to determine reliability of the communications channel on the network; and adjust the values of M and N according to said reliability of the communications channel*, as recited in claim 12.

The Examiner submits that Davis teaches these limitations in column 30, lines 65-67, column 32, lines 15-29, and column 31, lines 1-3. However, these passages describe

adjusting the send window size dependent on acknowledgments received, and on throughput measurements. There is nothing in Davis that describes determining reliability of a communications channel, as recited in claim 12, or adjusting send or receive window sizes according to such a determination.

For at least the reasons above, the rejection of claim 12 is unsupported by the cited art and removal thereof is respectfully requested.

Claims 36 and 56 include limitations similar to those discussed above regarding claim 12. Therefore, the arguments presented above apply with equal force to this claim, as well.

Claims 13-14, 37-38, and 57-58 include further limitations on the adjustment of send and receive window sizes dependent on the determined reliability of an communications channel. Since the cited art does not teach or suggest such a determination, it clearly does not teach or suggest these further limitations.

Second ground of rejection:

The Examiner rejected claims 4, 8-10, 28, 32-34, 48 and 52-54 as being unpatentable over Davis, Dreke and Black in view of Barker et al. (U.S. Patent 5,931,916) (hereinafter "Barker"). Appellants traverse this rejection for at least the following reasons. Different groups of claims are addressed under their respective subheadings.

Dependent claims 4, 28, and 48:

1. The cited art clearly fails to teach or suggest *wherein N is a positive even integer, and wherein M is equal to $N/2$* , as recited in claim 4.

The Examiner admits that Davis does not specifically teach this limitation. The examiner references Barker as teaching “a similar system of adjusting the window for the transmission of packets, wherein the receiver sends an acknowledgement after a certain number of messages in a sequence have been received (col. 6, lines 25-31, 63-66).” The Examiner further admits that Davis and Barker do not explicitly teach wherein M is equal to $N/2$, as recited in claim 4. The Examiner submits that since Barker teaches transmitting an acknowledgement after a certain number of messages, “it would have been obvious to one of ordinary skill for the receiver endnode to transmit an acknowledgement after other numbers of messages including $N/2$ messages.”

Appellants assert that The Examiner’s reason is not supported by any evidence of record and is thus found only in hindsight. **There is nothing in the evidence of record teaching or suggesting any reason (or benefit) for the number of messages to have the specific limitation recited in claim 4, wherein M is equal to $N/2$.** In fact, by stating that any other number of messages may be transmitted, the Examiner acknowledges that there is no particular benefit to any of the particular numbers.

For at least the reasons above, the rejection of claim 4 is unsupported by the cited art and removal thereof is respectfully requested.

Claims 28 and 48 include limitations similar to those discussed above regarding claim 4. Therefore, the arguments presented above apply with equal force to this claim, as well.

Dependent claims 8, 32, and 52:

Appellants traverse the rejection of these claims for at least the reasons presented above regarding the claims from which they depend.

Dependent claims 9, 10, 33, 34, 53, and 54:

1. ***The cited art clearly fails to teach or suggest wherein the other peer node is configured to: continue receiving the transmitted messages until the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received... wherein M is a positive integer less than N , as recited in claim 9.***

The Examiner submits that Davis teaches this limitation, citing column 6, lines 63-67 and column 73, lines 44-47. These passages describe acknowledging receipt of a plurality of packets, but do not teach or suggest anything about continuing to receive packets until a particular number of packets are received, as required by claim 9.

For at least the reasons above, the rejection of claim 9 is unsupported by the cited art and removal thereof is respectfully requested.

Claims 33 and 53 include limitations similar to those discussed above regarding claim 9. Therefore, the arguments presented above apply with equal force to this claim, as well.

Third ground of rejection:

The Examiner rejected claims 16 and 17 as being unpatentable over Davis, Dreke and Black in view of Ivanoff (U.S. Patent 5,517,622). Appellants traverse this rejection for at least the reasons presented above regarding the claim from which these claims depend.

Fourth ground of rejection:

The Examiner rejected claims 19, 20, 41, 42, 61 and 62 as being unpatentable over Davis, Dreke and Black in view of Antur et al. (U.S. Patent 6,212,558) (hereinafter “Antur”). Appellants traverse this rejection for at least the reasons presented above regarding the claims from which these claims depend.

Fifth ground of rejection:

The Examiner rejected claims 23, 24, 44 and 64 as being unpatentable over Davis, Dreke and Black in view of Zhu et al. (U.S. Patent 5,768,557) (hereinafter “Zhu”). Appellants note that U.S. Patent No. 5,768,557 is not a patent issued to an inventor “Zhu” and assert that this patent does not teach the limitations recited in claims 23, 24, 44, and 64. For at least these reasons, the Examiner has failed to establish a *prima facie* obviousness of the claimed invention.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-64 was erroneous, and reversal of his decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-07400/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Appellants

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: August 25, 2008

VIII. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A peer computing system, comprising:

a plurality of peer nodes operable to couple to a network, wherein each of the plurality of peer nodes comprises one or more network interfaces, wherein each network interface is configured to communicate over the network in accordance with at least one of one or more network transport protocols;

wherein the plurality of peer nodes is configured to implement a peer-to-peer environment on the network according to a peer-to-peer platform comprising one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and share content in the peer-to-peer environment, wherein to discover comprises obtaining an address for each discovered peer node;

wherein one of the plurality of peer nodes is configured to:

establish a communications channel between a network interface of the peer node and a network interface of another of the plurality of peer nodes;

transmit messages to the other peer node over the communications channel;

receive acknowledgement that one or more of the transmitted messages have been received by the other peer node; and

retransmit messages not acknowledged as received by the other peer node to the other peer node on the communications channel;

wherein said establishing, said transmitting, said receiving, and said retransmitting are performed according to at least one of the one or more peer-to-peer platform protocols and wherein said peer-to-peer platform protocols are distinct from the at least one network transport protocols.

2. The peer computing system as recited in claim 1, wherein, to transmit messages to the other peer node over the communications channel, the peer node is further configured to:

generate the messages;

buffer the messages, and after a window of N messages has been buffered, transmit the N messages to the other peer node over the communications channel, wherein N is an integer greater than one.

3. The peer computing system as recited in claim 2, wherein the other peer node is configured to receive the transmitted messages, and after receiving M messages, transmit the acknowledgement to the peer node indicating that the M messages have been received, wherein M is a positive integer less than or equal to N.

4. The peer computing system as recited in claim 3, wherein N is a positive even integer, and wherein M is equal to $N / 2$.

5. The peer computing system as recited in claim 3, wherein M is less than N.

6. The peer computing system as recited in claim 5, wherein, to receive acknowledgement that one or more of the transmitted messages have been received by the other peer node, the peer node is further configured to receive the acknowledgement indicating that the M messages have been received, and wherein the peer node is further configured to:

shift the window in the buffer by M messages; and

transmit the messages in the shifted window to the other peer node over the communications channel.

7. The peer computing system as recited in claim 6, wherein the shifted window includes one or more messages previously transmitted to the other peer node and one or more messages not previously transmitted to the other peer node.

8. The peer computing system as recited in claim 2, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the other peer node is configured to:

receive the transmitted messages; and

after receiving the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers, transmit the acknowledgement to the peer node indicating that the first M messages have been received, wherein M is a positive integer less than N.

9. The peer computing system as recited in claim 2, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the other peer node is configured to:

continue receiving the transmitted messages until the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received or a timeout limit from the time of initial receipt of one of the sequence of N transmitted messages is exceeded, wherein M is a positive integer less than N;

if the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received, transmit the acknowledgement to the peer node indicating that a count of messages received in continuous sequence from a first message in the sequence of N transmitted messages is M; and

if the timeout limit is exceeded before the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received, transmit the acknowledgement to the peer node indicating the count of messages received in continuous sequence from the first message in the sequence of N transmitted messages, wherein the count of messages received in continuous sequence is less than M.

10. The peer computing system as recited in claim 9, wherein, to receive acknowledgement that one or more of the transmitted messages have been received by the other peer node, the peer node is further configured to receive the acknowledgement indicating that the messages have been received, and wherein the peer node is further configured to:

shift the window in the buffer by the count of messages received in continuous sequence; and

transmit the messages in the shifted window to the other peer node over the communications channel.

11. The peer computing system as recited in claim 1, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node.

12. The peer computing system as recited in claim 3, wherein the peer node and the other peer node are further configured to:

monitor reception and retransmission of the messages to determine reliability of the communications channel on the network; and

adjust the values of M and N according to said reliability of the communications channel.

13. The peer computing system as recited in claim 12, wherein, to adjust the values of M and N, the peer node and the other peer node are further configured to lower the values of M and N if said reliability of the communications channel is poor.

14. The peer computing system as recited in claim 12, wherein, to adjust the values of M and N, the peer node and the other peer node are further configured to raise the values of M and N if said reliability of the communications channel is good.

15. The peer computing system as recited in claim 1, wherein the other peer node is configured to:

transmit other messages to the peer node over the communications channel;

receive acknowledgement that one or more of the transmitted other messages have been received by the peer node; and

retransmit messages not acknowledged as received by the peer node to the peer node on the communications channel.

16. The peer computing system as recited in claim 1, wherein the peer node comprises an instance of a pipe service executable within the peer node to establish the communications channel, transmit the messages to the other peer node, receive the acknowledgement, and retransmit the messages not acknowledged as received.

17. The peer computing system as recited in claim 16, wherein the other peer node comprises another instance of the pipe service executable within the other peer node to receive the transmitted messages and transmit the acknowledgement to the peer node.

18. The peer computing system as recited in claim 1, wherein the communications channel passes through one or more relay peers, wherein the one or more relay peers are configured to receive the transmitted messages from the peer node and forward the messages to the other peer node.

19. The peer computing system as recited in claim 1, wherein the communications channel passes through one or more firewalls.

20. The peer computing system as recited in claim 1, wherein the communications channel passes through one or more Network Address Translation (NAT) gateways.

21. The peer computing system as recited in claim 1, wherein one or more other of the plurality of peer nodes are configured to connect to the communications channel, wherein the peer node is further configured to:

transmit messages to the one or more other peer nodes over the communications channel;

receive acknowledgements that one or more of the transmitted messages have been received by the one or more other peer nodes; and

retransmit messages not acknowledged as received by the one or more other peer nodes to the one or more other peer node on the communications channel.

22. The peer computing system as recited in claim 1, wherein the peer node is further configured to compare elapsed time since the messages were transmitted to a timeout limit and, if the elapsed time exceeds the timeout limit, retransmit the messages to the other peer node over the communications channel.

23. The peer computing system as recited in claim 1, wherein the peer node is further configured to:

receive a request specifying one or more previously transmitted messages for retransmission by the peer node; and

retransmit the specified one or more messages to the other peer node on the communications channel in response to the request.

24. The peer computing system as recited in claim 23, wherein the request specifies a sequence number for each of the one or more specified messages, wherein the sequence numbers are for use in ordering the received messages on the other peer node.

25. A method for providing reliable connections between peer nodes coupled to a peer-to-peer network, the method comprising:

a plurality of peer nodes coupled to the network implementing a peer-to-peer environment on the network according to a peer-to-peer platform comprising one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each

other, and share content in the peer-to-peer environment, wherein to discover comprises obtaining an address for each discovered peer node;

establishing a communications channel between a network interface of one of the plurality of peer nodes and a network interface of another of the plurality of peer nodes;

the peer node transmitting messages to the other peer node over the communications channel;

the peer node receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node; and

the peer node retransmitting messages not acknowledged as received by the other peer node to the other peer node on the communications channel;

wherein said establishing, said transmitting, said receiving, and said retransmitting are performed according to at least one of the one or more peer-to-peer platform protocols, and wherein the at least one of the one or more peer-to-peer platform protocols is distinct from any underlying network transport protocols.

26. The method as recited in claim 25, wherein, in said transmitting messages to the other peer node over the communications channel, the method further comprises:

generating the messages;

buffering the messages, and after a window of N messages has been buffered, transmitting the N messages to the other peer node over the communications channel, wherein N is an integer greater than one.

27. The method as recited in claim 26, further comprising the other peer node receiving the transmitted messages, and after receiving M messages, transmitting the acknowledgement to the peer node indicating that the M messages have been received, wherein M is a positive integer less than or equal to N.

28. The method as recited in claim 27, wherein N is a positive even integer, and wherein M is equal to $N / 2$.

29. The method as recited in claim 27, wherein M is less than N.

30. The method as recited in claim 29, wherein, in said receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, the method further comprises the peer node receiving the acknowledgement indicating that the M messages have been received, and wherein the method further comprises:

shifting the window in the buffer by M messages; and

transmitting the messages in the shifted window to the other peer node over the communications channel.

31. The method as recited in claim 30, wherein the shifted window includes one or more messages previously transmitted to the other peer node and one or more messages not previously transmitted to the other peer node.

32. The method as recited in claim 26, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the method further comprises:

the other peer node receiving the transmitted messages; and

after receiving the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers, the other peer node transmitting the acknowledgement to the peer node indicating that the first M messages have been received, wherein M is a positive integer less than N.

33. The method as recited in claim 26, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the method further comprises:

the other peer node continuing to receive the transmitted messages until the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received or a timeout limit from the time of initial receipt of one of the sequence of N transmitted messages is exceeded, wherein M is a positive integer less than N;

if the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received, the other peer node transmitting the acknowledgement to the peer node indicating that a count of messages received in continuous sequence from a first message in the sequence of N transmitted messages is M; and

if the timeout limit is exceeded before the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers are received, the other peer node transmitting the acknowledgement to the peer node indicating the count of messages received in continuous sequence from the first message in the sequence of N transmitted messages, wherein the count of messages received in continuous sequence is less than M.

34. The method as recited in claim 33, wherein, in said receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, the method further comprises the peer

node receiving the acknowledgement indicating that the messages have been received, and wherein the method further comprises:

shifting the window in the buffer by the count of messages received in continuous sequence; and

transmitting the messages in the shifted window to the other peer node over the communications channel.

35. The method as recited in claim 25, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node.

36. The method as recited in claim 27, further comprising:

monitoring reception and retransmission of the messages to determine reliability of the communications channel on the network; and

adjusting the values of M and N according to said reliability of the communications channel.

37. The method as recited in claim 36, wherein, in said adjusting the values of M and N, the method further comprises lowering the values of M and N if said reliability of the communications channel is poor.

38. The method as recited in claim 36, wherein, in said adjusting the values of M and N, the method further comprises raising the values of M and N if said reliability of the communications channel is good.

39. The method as recited in claim 25, further comprising:

the other peer node transmitting other messages to the peer node over the communications channel;

the other peer node receiving acknowledgement that one or more of the transmitted other messages have been received by the peer node; and

the other peer node retransmitting messages not acknowledged as received by the peer node to the peer node on the communications channel.

40. The method as recited in claim 25, wherein the communications channel passes through a relay peer, the method further comprising the relay peer receiving the transmitted messages from the peer node and forwarding the messages to the other peer node.

41. The method as recited in claim 25, wherein the communications channel passes through one or more firewalls.

42. The method as recited in claim 25, wherein the communications channel passes through one or more Network Address Translation (NAT) gateways.

43. The method as recited in claim 25, further comprising the peer node comparing elapsed time since the messages were transmitted to a timeout limit and, if the elapsed time exceeds the timeout limit, retransmitting the messages to the other peer node over the communications channel.

44. The method as recited in claim 25, further comprising:

the peer node receiving a request specifying one or more previously transmitted messages for retransmission by the peer node; and

the peer node retransmitting the specified one or more messages to the other peer node on the communications channel in response to the request.

45. A computer-readable storage medium, comprising software instructions executable on a peer node to implement:

the peer node establishing a communications channel between a network interface of the peer node and a network interface of another peer node of a plurality of peer nodes coupled to a network implementing a peer-to-peer environment on the network according to a peer-to-peer platform comprising one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and share content in the peer-to-peer environment, wherein to discover comprises obtaining an address for each discovered peer node;

the peer node transmitting messages to the other peer node over the communications channel;

the peer node receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node; and

the peer node retransmitting messages not acknowledged as received by the other peer node to the other peer node on the communications channel;

wherein said establishing, said transmitting, said receiving, and said retransmitting are performed according to at least one of the one or more peer-to-peer platform protocols, and wherein the at least one of the one or more peer-to-peer platform protocols is distinct from any underlying network transport protocols.

46. The storage medium as recited in claim 45, wherein, in said transmitting messages to the other peer node over the communications channel, the software instructions are further executable to implement:

generating the messages;

buffering the messages, and after a window of N messages has been buffered, transmitting the N messages to the other peer node over the communications channel, wherein N is an integer greater than one.

47. The storage medium as recited in claim 46, wherein the software instructions are further executable to implement the peer node receiving the acknowledgement from the other peer node indicating that the M messages have been received by the other peer node, wherein M is a positive integer less than or equal to N.

48. The storage medium as recited in claim 47, wherein N is a positive even integer, and wherein M is equal to $N / 2$.

49. The storage medium as recited in claim 47, wherein M is less than N.

50. The storage medium as recited in claim 49, wherein, in said receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, the software instructions are further executable to implement:

shifting the window in the buffer by M messages; and

transmitting the messages in the shifted window to the other peer node over the communications channel.

51. The storage medium as recited in claim 50, wherein the shifted window includes one or more messages previously transmitted to the other peer node and one or more messages not previously transmitted to the other peer node.

52. The storage medium as recited in claim 46, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the software instructions are further executable to implement:

the peer node receiving acknowledgement from the other peer node indicating that the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers have been received by the other peer node, wherein M is a positive integer less than N.

53. The storage medium as recited in claim 46, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node, and wherein the software instructions are further executable to implement:

the peer node receiving an acknowledgement from the other peer node indicating:

the other peer node received M messages in continuous sequence from a first message in the sequence of N transmitted message if the other peer node continued to receive the transmitted messages until the first M messages in the sequence of N transmitted messages were received; or

the other peer node received less than M messages in continuous sequence from the first message in the sequence of N transmitted messages if a timeout limit from initial receipt of one of the sequence of N transmitted message was exceeded before the first M messages in the sequence of N transmitted messages as indicated by the sequence numbers was received by the other peer node.

54. The storage medium as recited in claim 53, wherein, in said receiving acknowledgement that one or more of the transmitted messages have been received by the other peer node, the software instructions are further executable to implement the peer node receiving the acknowledgement indicating that the messages have been received, and wherein the software instructions are further executable to implement:

shifting the window in the buffer by the count of messages received in continuous sequence; and

transmitting the messages in the shifted window to the other peer node over the communications channel.

55. The storage medium as recited in claim 45, wherein each of the messages includes a sequence number for use in ordering the received messages on the other peer node.

56. The storage medium as recited in claim 47, wherein the software instructions are further executable to implement:

monitoring reception and retransmission of the messages to determine reliability of the communications channel on the network; and

adjusting the values of M and N according to said reliability of the communications channel.

57. The storage medium as recited in claim 56, wherein, in said adjusting the values of M and N, the software instructions are further executable to implement lowering the values of M and N if said reliability of the communications channel is poor.

58. The storage medium as recited in claim 56, wherein, in said adjusting the values of M and N, the software instructions are further executable to implement raising the values of M and N if said reliability of the communications channel is good.

59. The storage medium as recited in claim 45, wherein the software instructions are further executable to implement:

the peer node receiving other messages from the other peer node over the communications channel;

the peer node sending acknowledgement to the other peer node that one or more of the transmitted other messages have been received by the peer node; and

the peer node receiving retransmitted messages not acknowledged as received by the peer node on the communications channel.

60. The storage medium as recited in claim 45, wherein the software instructions are further executable to implement:

configuring the peer node as a relay peer, wherein a communications channel between a third peer node of the plurality of peer nodes and the other peer node passes through the peer node;

the peer node receiving messages transmitted from the third peer node to the other peer node; and

forwarding the messages to the other peer node.

61. The storage medium as recited in claim 45, wherein the communications channel passes through one or more firewalls.

62. The storage medium as recited in claim 45, wherein the communications channel passes through one or more Network Address Translation (NAT) gateways.

63. The storage medium as recited in claim 45, wherein the software instructions are further executable to implement the peer node comparing elapsed time since the messages were transmitted to a timeout limit and, if the elapsed time exceeds the timeout limit, retransmitting the messages to the other peer node over the communications channel.

64. The storage medium as recited in claim 45, wherein the software instructions are further executable to implement:

the peer node receiving a request specifying one or more of the transmitted messages for retransmission by the peer node; and

the peer node retransmitting the specified one or more messages to the other peer node on the communications channel in response to the request.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.